

# Toward Sequentializing Overparallelized Protocol Code

Sung-Shik T.Q. Jongmans

Farhad Arbab

Centrum Wiskunde & Informatica  
Amsterdam, Netherlands

[jongmans,farhad]@cwi.nl

In our ongoing work, we use constraint automata to compile protocol specifications expressed as Reo connectors into efficient executable code, e.g., in C. We have by now studied this automata based compilation approach rather well, and have devised effective solutions to some of its problems. Because our approach is based on constraint automata, the approach, its problems, and our solutions are in fact useful and relevant well beyond the specific case of compiling Reo. In this short paper, we identify and analyze two such rather unexpected problems.

## Introduction

A promising application domain for coordination languages is programming protocols among threads in multicore programs: coordination languages typically provide high-level constructs and abstractions that more easily compose into correct—with respect to a programmer’s intentions—protocol specifications than do low-level synchronization constructs provided by conventional languages (e.g., locks, semaphores). In fact, not only do coordination languages simplify programming protocols, but their high-level constructs and abstractions also leave more room for compilers to perform novel optimizations in mapping protocol specifications to lower-level instructions that implement them. A crucial step toward adoption of coordination languages for multicore programming is the development of such compilers: programmers need tools to generate efficient code from high-level protocol specifications.

In ongoing work, we develop compiler technology for the graphical coordination language Reo [1]. Reo facilitates compositional construction of protocol specifications manifested as *connectors*: channel-based mediums through which threads can communicate with each other. Figure 1 shows some example connectors, each linked to four *computation threads*, in their usual graphical syntax. Briefly, a connector consists of one or more *channels*, through which data items flow, and a number of *nodes*, on which channel ends coincide. In Figure 1, we distinguish the *boundary nodes* of a connector (to which computation threads are linked) from its *internal nodes* (used only for internally routing data) by shading the

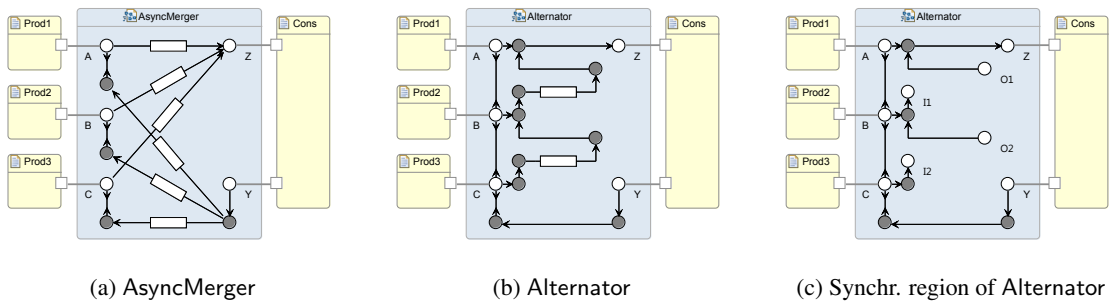


Figure 1: Example connectors

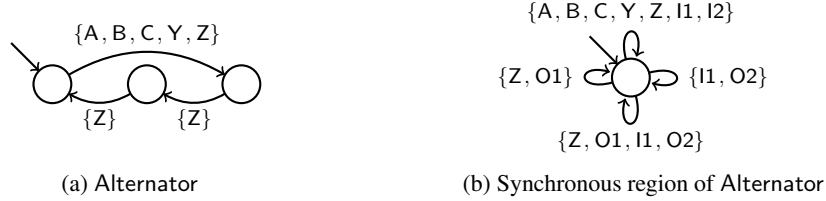


Figure 2: Example constraint automata (irrelevant details of transition labels omitted)

internal nodes. The connectors in Figure 1 contain three different channel classes, including standard synchronous channels (normal arrows) and asynchronous channels with a buffer of capacity 1 (arrows decorated with a white rectangle, which represents a buffer). Through connector *composition* (the act of gluing connectors together on their shared nodes), programmers can construct arbitrarily complex connectors. As Reo supports both synchronous and asynchronous channels, connector composition enables mixing synchronous and asynchronous communication within the same protocol specification.

Figure 1a shows a connector, AsyncMerger, for a protocol among  $k = 3$  producers and one consumer. We compared the code generated by our Reo-to-C compiler [15] with hand-crafted code written by a competent C programmer using Pthreads, investigating the time required for communicating a data item from a producer to the consumer as a function of the number of producers  $4 \leq k \leq 512$ .

The results looked excellent: the code generated by our compiler outperforms the hand-crafted code and scales well [14]. Encouraged by this outcome, we expected to reproduce these results for the producers–consumer protocol specified by the Alternator connector in Figure 1b.<sup>1</sup> The results disappointed us: for small  $k$ , the code of Alternator runs significantly slower than that of AsyncMerger, while for large  $k$ , the compiler times out (i.e., after five minutes, we manually aborted the compilation process).

In this short paper, we identify two “unexpected” problems of our current compilation approach (which manifest in Alternator): exponential explosion at compile-time and overparallelization at run-time. These problems are in fact unfortunate side effects of another optimization step in our compilation process that we thought we had well studied. After an analysis, we propose a first solution that works in some—but not all—problematic cases; we leave a comprehensive solution for future work and consider the identification and analysis of the two problems the main contribution of this short paper.

## Problem Analysis and a First Solution

Our Reo-to-C compiler generates code for Reo connectors based on their *constraint automaton* (CA) semantics [3]. Constraint automata are a general formalism for modeling systems, better suited for data-aware modeling of Reo connectors and, in particular, their composition (which supports multiparty and transitive synchronization) than classical automata or traditional process calculi. Figure 2a shows examples. For Reo, a CA specifies *when* during execution of a connector *which* data items flow *where*. Structurally, every CA consists of finite sets of states and transitions. A *product operator* on CA, which preserves CA-bisimilarity [3], models connector composition: to obtain the “big” CA for a whole connector, one can compute the product of the “small” CA for its constituent nodes and channels. Afterward, one can abstract away internal nodes with a *hide operator* on CA [3], which—importantly—also eliminates silent transitions involving only internal nodes in a semantics-preserving way.

<sup>1</sup>In the AsyncMerger protocol, the consumer receives productions in arbitrary order. In contrast, in the Alternator protocol, the consumer receives data “from top to bottom” (and to achieve this, the producers collectively synchronize before sending).

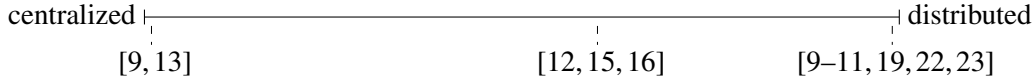


Figure 3: Connector implementation spectrum

Although motivated by our work on Reo, our compiler really operates primarily at the level of Reo’s CA semantics. In that sense, “Reo-to-C compiler” is a misnomer. A better name would be “CA-to-C compiler”: we use Reo, with its graphical, channel-based abstractions, just as *a*—not *the*—programmer-friendly syntax for exposing CA-based protocol programming. Different syntax alternatives for CA may work equally well or yield perhaps even more user-friendly languages. For instance, we know how to translate UML sequence/activity diagrams and BPMN to CA [2, 8, 18]. Another interesting potential syntax are algebras of Bliudze and Sifakis [5], originally developed in the context of BIP [4], which have a straightforward interpretation in terms of CA. Due to their generality, CA can thus serve as an intermediate language (transparent to programmers) for compiling specifications in many different languages and models of concurrency by reusing the core of our compiler. This makes the development of this compiler and its optimizations relevant beyond Reo.

Two opposite CA-based approaches to implementing a connector Conn exist. In the *distributed approach*, the compiler first finds a small CA for every channel and every node that Conn consists of and afterward generates a piece of sequential code for each of those small CA. At run-time, every piece of sequential code has its own thread, henceforth referred to as *protocol threads*, and a distributed algorithm among those threads ensures their proper synchronization. In the *centralized approach*, after finding a collection of small CA, the compiler forms the product of all those CA to get a big CA for Conn, abstracts away all internal nodes, and finally generates one piece of sequential code for that big CA. For CA-based implementations, these two approaches constitute the two ends of the *connector implementation spectrum* in Figure 3: the further we get to the right end of the spectrum, the more parallelism a connector implementation exhibits. (For completeness, Figure 3 contains also references to Reo connector implementation approaches based on other formalisms—in particular, *connector coloring* and *coordination constraints* [9–11, 19–23]—which work not exactly the same as just described for CA.)

Neither the distributed approach nor the centralized approach is satisfactory. For instance, the distributed approach suffers from high latency at run-time (because the distributed algorithm required for synchronizing the parallel protocol threads is expensive). The centralized approach, in contrast, achieves low latency, but it suffers from state space explosion at compile-time (because a big CA for a whole connector may have a number of states exponential in the number of its constituent channels) and oversequentialization at run-time (because simulating a big CA with one thread serializes transitions that could have fired in parallel). To solve these problems (i.e., strike a balance between run-time latency and parallelism), we extensively studied a *middle ground approach* roughly in the center of the connector implementation spectrum. In this approach, the compiler splits a connector into  $m_1$  *asynchronous regions* of purely asynchronous communication (e.g., each of the buffered channels in Figure 1) and  $m_2$  *synchronous regions* of synchronous communication.<sup>2</sup> The compiler subsequently forms products on a per-region basis, resulting in  $m_1 + m_2$  “medium” CA, and generates a piece of sequential code for each of them. At run-time, every generated piece of code has its own thread, as in the distributed approach, but the distributed algorithm required for synchronizing those protocol threads has substantially lower costs. Moreover, the middle ground approach mitigates state space explosion and oversequentialization. For these advantages, we moved our compiler from the centralized approach to the middle ground approach.

<sup>2</sup>Splitting into regions occurs at the level of small CA, without knowledge of the input connector [12, 16].

Unfortunately and unexpectedly, although the middle ground approach works well for AsyncMerger, it fails for Alternator. We analyze why, as follows. First, Figure 1c shows the single synchronous region of Alternator. Because nodes  $l1$ ,  $l2$ ,  $O1$ , and  $O2$  lie on the boundary of this region, the compiler cannot abstract those nodes away. Next, Figure 2b shows the medium CA for this region. Its  $\{A, B, C, Y, Z, l1, l2\}$ -transition and its  $\{Z, O1\}$ -transition correspond to the  $\{A, B, C, Y, Z\}$ -transition and the two  $\{Z\}$ -transitions of the big CA in Figure 2a. The  $\{l1, O2\}$ -transition of the medium CA models an internal execution step—abstracted away in the big CA—in which a data item flows from the bottom buffer into the top buffer. Finally, the  $\{Z, O1, l1, O2\}$ -transition of the medium CA models an execution step in which its  $\{Z, O1\}$ -transition and its  $\{l1, O2\}$ -transition fire simultaneously by true concurrency.

Now, imagine a generalization of Alternator from three producers to  $k$  producers (by replicating parts of Alternator in Figure 1b in the obvious way). Such a connector has  $k - 1$  buffers. Consequently, the medium CA for its single synchronous region has  $k - 2$  transitions (among others), each of which models an internal execution steps where a data item flows from one buffer to the buffer directly above it. Because any subset of those transitions may fire simultaneously by true concurrency, the medium CA has roughly  $2^{k-2}$  transitions. The medium CA for Alternator with 512 producers consequently has over  $10^{153}$  transitions—approximately  $10^{73}$  times the estimated number of hydrogen atoms in the observable universe—such that merely representing this CA in memory is already problematic (let alone compositionally computing it). Thus, transition relation explosion at compile-time is a serious problem.

Now, suppose that we manage to successfully compile Alternator for a sufficiently small number of  $\ell$  producers. At run-time, we have  $\ell$  parallel protocol threads: one for Alternator’s synchronous region and one for each of its  $\ell - 1$  asynchronous regions. But despite this parallel implementation, the big CA of Alternator in Figure 2a (for  $\ell = 3$ ) implies that Alternator in fact behaves sequentially. In other words, we use parallelism—and incur the overhead that parallelism involves—to implement intrinsically sequential behavior. Thus, overparallelization at run-time is another serious problem.

Interestingly, the centralized approach, which our compiler used to apply, does not suffer from transition relation explosion or overparallelization for a number of reasons. First, overparallelization is trivially not a problem, because the centralized approach involves only one sequential protocol thread.

The second reason relates to the fact that enabledness of transitions in Alternator’s synchronous region depends on the (non)emptiness of the buffers in its  $k$  asynchronous regions: many transitions are in fact permanently disabled. For instance, every “true-concurrency-transition” composed of  $3 \leq x \leq k - 1$  transitions labeled with  $\{li, Oi + 1\}$  (for some  $i$ ), where data items flow upward through  $x$  consecutive buffers, never fires: by Reo’s semantics, the  $x - 2$  middle buffers cannot become empty and full again in the same transition, which would happen if this true-concurrency-transition were to fire. A compiler can eliminate such permanently disabled transitions—and thereby mitigate transition relation explosion—by forming the product of all medium CA for Alternator’s synchronous and asynchronous regions (in a particular order), effectively computing one big CA. Exactly this happens in the centralized approach.

The third reason relates to abstraction of internal nodes and transitions. In the middle ground approach, nodes shared between different regions do not count as internal nodes; they are boundary nodes and the compiler cannot abstract them away. In contrast, in the centralized approach, all those boundary nodes between regions become internal nodes, which the compiler can abstract away. Consequently, the compiler can eliminate more silent transitions involving only internal nodes—and thereby further mitigate transition relation explosion—by applying the hide operator.

Having moved our compiler from the centralized approach to the middle ground approach to avoid state space explosion and oversequentialization, now, we must find solutions for the unfortunate side effects of this move: transition relation explosion and overparallelization.

Our first solution is to, at compile-time, merge every asynchronous region AR that shares nodes

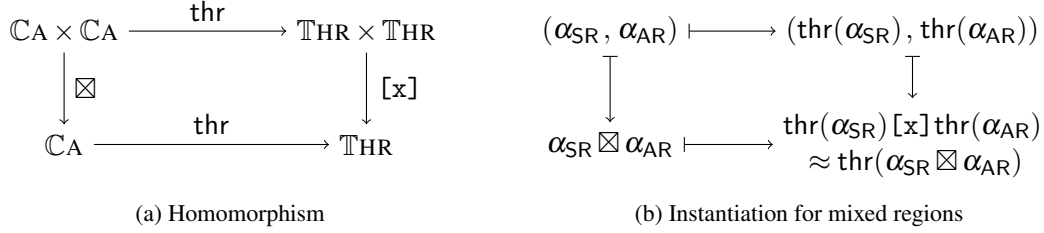


Figure 4: Justification of mixed regions, where  $\mathbb{CA}$  denotes the set of all CA,  $\mathbb{THR}$  denotes the set of all protocol threads,  $\text{thr}$  denotes a translation from CA to protocol threads (i.e., actual code generation),  $\boxtimes$  denotes the product operator on CA,  $[x]$  denotes parallel composition of protocol threads synchronized by a distributed algorithm [12, 16], and  $\approx$  denotes observational equivalence of protocol threads.

with only one synchronous region SR (i.e., AR is neither connected to another region nor linked to a computation thread) into SR. Doing so results in a *mixed region*. Computation of mixed regions is semantics-preserving by the associativity and commutativity of the product operator on CA [3]: if  $\alpha_{\text{SR}}$ ,  $\alpha_{\text{AR}}$ , and  $\alpha_{\text{other}}$  denote the CA for SR, AR, and the other regions, the compiler can always change the bracketing of a product term over those CA to a form in which  $\alpha_{\text{SR}}$  and  $\alpha_{\text{AR}}$  are the operands of the same product operator. The compiler can subsequently decide either to actually form that product (thus computing the CA of a mixed region) or leave  $\alpha_{\text{SR}}$  and  $\alpha_{\text{AR}}$  as separate CA. In the former case, at run-time, the protocol thread for the resulting product participates as one entity in the distributed algorithm for synchronizing protocol threads; in the latter case, both the protocol thread for  $\alpha_{\text{SR}}$  and the protocol thread for  $\alpha_{\text{AR}}$  participate in this algorithm. Semantically, these implementations are indistinguishable. More formally, the diagram in Figure 4 commutes.

Intuitively, forming mixed regions mitigates transition relation explosion at compile-time because (i) the compiler essentially computes a bigger product (which may eliminate permanently disabled transitions) and (ii) the compiler can abstract away more internal nodes (which may eliminate more silent transitions involving only internal nodes), namely all those shared between SR and AR. Overparallelization at run-time is mitigated because every asynchronous region connected only to SR *must* interact with SR in each of its transitions; it can never fire a transition independently of SR. Running such an asynchronous region in its own protocol thread would therefore never result in useful parallelism.

If we apply this first solution to Alternator, the compiler merges all asynchronous regions into Alternator’s single synchronous region. This results in a single mixed region spanning the whole connector. In this case, thus, the compiler reduces the middle ground approach back to centralized approach.

Although formulated generally in terms of regions, we know of cases of overparallelization that our first solution fails to mitigate. For instance, although the Sequencer connector has intrinsically sequential behavior [1], each of its asynchronous regions has connections to two—not one—synchronous regions. We are thinking of generalizing our first solution to capture also this and similar cases, although we are not convinced yet that such a generalization exists; perhaps we need a rather different kind of rule.

Generally, two behaviorally equivalent but structurally different connectors may yield different pieces of code with different performance. Figure 5 shows three behaviorally equivalent connectors demonstrating that this applies also to the problems identified in this short paper. (To see this, note that because the connector in Figure 5c includes Alternator, it suffers from the same problems as Alternator). Consequently, another solution for these problems may be to structurally manipulate connectors (or the sets of small CA they behave as) before splitting them into regions. Although we conjecture that such manipulation not always solves our problems, we may identify a class of connectors for which it does.

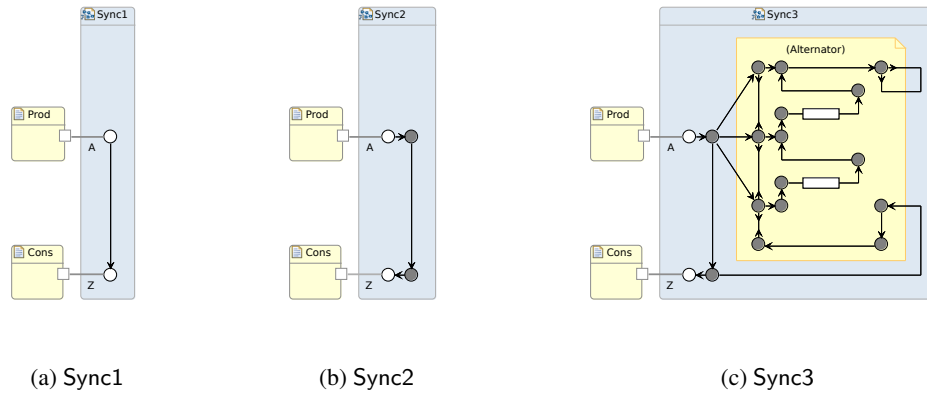


Figure 5: Behaviorally equivalent connectors, each of which models a standard synchronous channel.

Finally, at least transition relation explosion may be mitigated by improving our way of dealing with parametrization. In the Alternator case, for instance, our current approach to (static) parametrization problematically requires the compiler to compute the CA for the whole  $k$ -sized region, given  $k$  producers. A better approach to (static or dynamic) parametrization may enable direct generation of code for  $k$  based on the CA for a 2-sized region without ever computing the CA for the whole  $k$ -sized region.

## Conclusion

We introduced two problems—transition relation explosion and overparallelization—with our current compilation approach for Reo. Intuitively, these problems can be regarded as the flip side of oversequentialization, and its accompanying plague of state space explosion. Although our first solution works in some cases, a comprehensive solution (including a better understanding of all the cases that this solution should cover), needs to be further developed. Essentially, we aim at finding the optimal position in the connector implementation spectrum in Figure 3 that perfectly balances parallelism and sequentiality.

Although encountered by us in the context of Reo, mitigating overparallelization seems a generally interesting problem. For instance, specifying a system as many parallel processes may feel natural to a system architect, but implementing each of those processes as a thread may give poor performance. By studying this problem in terms of CA, which are related to process languages with multiparty synchronization [17], we hope to gain new insight and advance compilation technology in areas other than Reo too. As another example, automatically partitioning BIP interaction specifications for generating optimal distributed implementations is still an open problem [6, 7]. Further studies may clarify the extent to which the correspondence between BIP interactions and CA can be leveraged by reusing results on CA.

## References

- [1] Farhad Arbab (2011): *Puff, The Magic Protocol*. In: *Talcott Festschrift, LNCS 7000*, Springer, pp. 169–206, doi:10.1007/978-3-642-24933-4\_9.
- [2] Farhad Arbab, Natallia Kokash & Sun Meng (2008): *Towards Using Reo for Compliance-Aware Business Process Modeling*. In: *Proceedings of ISoLA 2008, CCIS 17*, Springer, pp. 108–123, doi:10.1007/978-3-540-88479-8\_9.

- [3] Christel Baier, Marjan Sirjani, Farhad Arbab & Jan Rutten (2006): *Modeling component connectors in Reo by constraint automata*. *SCP* 61(2), pp. 75–113, doi:10.1016/j.scico.2005.10.008.
- [4] Ananda Basu, Marius Bozga & Joseph Sifakis (2006): *Modeling Heterogeneous Real-time Components in BIP*. In: *Proceedings of SEFM 2006*, IEEE, pp. 3–12, doi:10.1109/SEFM.2006.27.
- [5] Simon Bliudze & Joseph Sifakis (2010): *Causal semantics for the algebra of connectors*. *FMSD* 36(2), pp. 167–194, doi:10.1007/s10703-010-0091-z.
- [6] Borzoo Bonakdarpour, Marius Bozga, Mohamad Jaber, Jean Quilbeuf & Joseph Sifakis (2012): *A framework for automated distributed implementation of component-based models*. *Distributed Computing* 25(5), pp. 383–409, doi:10.1007/s00446-012-0168-6.
- [7] Borzoo Bonakdarpour, Marius Bozga & Jean Quilbeuf (in press): *Model-based implementation of distributed systems with priorities*. *DAES*, doi:10.1007/s10617-012-9091-0.
- [8] Behnaz Changizi, Natallia Kokash & Farhad Arbab (2010): *A Unified Toolset for Business Process Model Formalization*. In: *Preproceedings of FESCA 2010*, pp. 147–156.
- [9] Dave Clarke, David Costa & Farhad Arbab (2007): *Connector colouring I: Synchronisation and context dependency*. *SCP* 66(3), pp. 205–225, doi:10.1016/j.scico.2007.01.009.
- [10] Dave Clarke & José Proença (2012): *Partial Connector Colouring*. In: *Proceedings of COORDINATION 2012*, LNCS 7274, Springer, pp. 59–73, doi:10.1007/978-3-642-30829-1\_5.
- [11] Dave Clarke, José Proença, Alexander Lazovik & Farhad Arbab (2011): *Channel-based coordination via constraint satisfaction*. *SCP* 76(8), pp. 681–710, doi:10.1016/j.scico.2010.05.004.
- [12] Sung-Shik Jongmans & Farhad Arbab (2013): *Global Consensus through Local Synchronization*. In: *Proceedings of FOCLASA 2013*, CCIS 393, Springer, pp. 174–188, doi:10.1007/978-3-642-45364-9\_15.
- [13] Sung-Shik Jongmans & Farhad Arbab (2013): *Modularizing and Specifying Protocols among Threads*. In: *Proceedings of PLACES 2012*, EPTCS 109, CoRR, pp. 34–45, doi:10.4204/EPTCS.109.6.
- [14] Sung-Shik Jongmans, Sean Halle & Farhad Arbab (2014): *Automata-based Optimization of Interaction Protocols for Scalable Multicore Platforms*. In: *Proceedings of COORDINATION 2014*, LNCS 8459, Springer, pp. 65–82, doi:10.1007/978-3-662-43376-8\_5.
- [15] Sung-Shik Jongmans, Sean Halle & Farhad Arbab (in press): *Reo: A Dataflow Inspired Language for Multi-core*. In: *Proceedings of DFM 2013*, IEEE.
- [16] Sung-Shik Jongmans, Francesco Santini & Farhad Arbab (2014): *Partially-Distributed Coordination with Reo*. In: *Proceedings of PDP 2014*, IEEE, pp. 697–706, doi:10.1109/PDP.2014.19.
- [17] Natallia Kokash, Christian Krause & Erik de Vink (2012): *Reo+mCRL2: A framework for model-checking dataflow in service compositions*. *FAC* 24(2), pp. 187–216, doi:10.1007/s00165-011-0191-6.
- [18] Sun Meng, Farhad Arbab & Christel Baier (2011): *Synthesis of Reo circuits from scenario-based interaction specifications*. *SCP* 76(8), pp. 651–680, doi:10.1016/j.scico.2010.03.002.
- [19] José Proença (2011): *Synchronous Coordination of Distributed Components*. Ph.D. thesis, Leiden University.
- [20] José Proença & Dave Clarke (2013): *Data Abstraction in Coordination Constraints*. In: *Proceedings of FOCLASA 2013*, CCIS 393, Springer, pp. 159–173, doi:10.1007/978-3-642-45364-9\_14.
- [21] José Proença & Dave Clarke (2013): *Interactive Interaction Constraints*. In: *Proceedings of COORDINATION 2013*, LNCS 7890, Springer, pp. 211–225, doi:10.1007/978-3-642-38493-6\_15.
- [22] José Proença, Dave Clarke, Erik de Vink & Farhad Arbab (2011): *Decoupled execution of synchronous coordination models via behavioural automata*. In: *Proceedings of FOCLASA 2011*, EPTCS 58, CoRR, pp. 65–79, doi:10.4204/EPTCS.58.5.
- [23] José Proença, Dave Clarke, Erik de Vink & Farhad Arbab (2012): *Dreams: a framework for distributed synchronous coordination*. In: *Proceedings of SAC 2012*, ACM, pp. 1510–1515, doi:10.1145/2245276.2232017.